

The significance of the integrated Sachs-Wolfe effect revisited

Tommaso Giannantonio^{1,2*}, Robert Crittenden³, Robert Nichol^{3,4}, Ashley J. Ross³

¹Ludwig-Maximilians-Universität München, Universitäts-Sternwarte München, Scheinerstr. 1, D-81679 München, Germany

²Excellence Cluster Universe, Technical University Munich, Boltzmannstr. 2, D-85748 Garching bei München, Germany

³Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, PO1 3FX, UK

⁴SEPnet, South East Physics Network

12 September 2012

ABSTRACT

We revisit the state of the integrated Sachs-Wolfe (ISW) effect measurements in light of newly available data and address criticisms about the measurements which have recently been raised. We update the data set previously assembled by Giannantonio et al. (2008b) to include new data releases for both the cosmic microwave background (CMB) and the large-scale structure (LSS) of the Universe. We find that our updated results are consistent with previous measurements. By fitting a single template amplitude, we now obtain a combined significance of the ISW detection at the 4.4σ level, which fluctuates by $\sim 0.4\sigma$ when alternative data cuts and analysis assumptions are considered. We also make new tests for systematic contaminations of the data, focusing in particular on the issues raised by Sawangwit et al. (2010). Amongst them, we address the rotation test, which aims at checking for possible systematics by correlating pairs of randomly rotated maps. We find results consistent with the expected data covariance, no evidence for enhanced correlation on any preferred axis of rotation, and therefore no indication of any additional systematic contamination. We publicly release the results, the covariance matrix, and the sky maps used to obtain them.

Key words: Cosmic background radiation; Large-scale structure of the Universe; Cosmology: observations.

1 INTRODUCTION

Observational evidence indicates the expansion of the Universe is accelerating at late times, which may be explained by a small cosmological constant, some negative-pressure dark energy fluid (Frieman et al. 2008), by modifications of the laws of gravity (see e.g. Clifton et al. 2012), or by some non-trivial distribution of the local large-scale structure (see e.g. Dunsby et al. 2010). Evidence for this acceleration is provided by multiple complementary probes, such as observations of distant Type Ia supernovae (Amanullah et al. 2010; Lampeitl et al. 2010a,b), cosmic microwave background (CMB) anisotropies (Komatsu et al. 2011), baryon acoustic oscillations (BAO) (Eisenstein et al. 2005; Percival et al. 2010), clusters of galaxies (Roza et al. 2010), and the integrated Sachs-Wolfe (ISW) effect (Giannantonio et al. 2008b; Ho et al. 2008).

We shall focus here on the latter, which consists of small secondary fluctuations in the CMB which are produced whenever gravitational potentials are evolving, as happens at late times in the case of the Universe undergoing a transition to a curvature- or dark energy-dominated phase (Sachs & Wolfe 1967). If we assume a flat universe as supported by the primary CMB data, then a de-

tection of the ISW represents a measurement of dark energy and its properties. Unfortunately, the amplitude of the ISW signal is small compared with the intrinsic CMB temperature anisotropies, contributing to the signal only at large scales. To overcome this, a technique was introduced to extract the ISW signal by cross-correlating the observed CMB with tracers of the local large-scale structure (LSS) of the Universe, such as wide-area galaxy catalogues (Crittenden & Turok 1996). As we will review below, this method has been used to successfully detect this signature of dark energy by many authors using several different LSS catalogues and the WMAP data of the CMB. More recently, multiple data sets were analysed jointly to maximise the extracted signal (Ho et al. 2008; Giannantonio et al. 2008b) (G08 herein), thus detecting the ISW signal at an overall significance of $\sim 4\sigma$, when fit a single amplitude.

However, some concerns have been raised about these detections, notably by Sawangwit et al. (2010); Francis & Peacock (2010a,b); Hernández-Monteagudo (2010); López-Corredoira et al. (2010). These concerns relate to three main areas: conflicting estimates of the statistical significance; searches based on new photometric data sets; and the possibility of larger than expected systematic contaminations.

The purpose of this paper is twofold. First we update our analysis to include the latest data of both the CMB and the large-scale

* E-mail: tommaso.giannantonio at usm.lmu.de

structure, using the 7-year WMAP maps and the latest available releases of the Sloan Digital Sky Survey (SDSS); we also publicly release the results of our analysis and our sky maps. Second, we evaluate the above-mentioned criticism and re-assess the overall state of the ISW measurements, focusing in particular on addressing the concerns by Sawangwit et al. (2010) (S10 herein).

The plan of the paper is as follows: after reviewing the analysis techniques and the current state of the ISW measurements in Section 2, we will describe our updated data set and the publicly released maps in Section 3. We then move on to a discussion of systematic uncertainties in Section 4, where we address a particular type of systematic test which has been discussed in S10 (the rotation test), finding results consistent with those expected given the covariance of the data. Further issues raised by S10 and other authors are addressed in Section 5, before we conclude in Section 6.

2 THE STATE OF THE ISW

2.1 Theory

The ISW effect (Sachs & Wolfe 1967) is a secondary source of temperature anisotropy which is produced whenever the gravitational potentials Φ and Ψ are evolving in time, generating temperature anisotropies of the form

$$\Theta_{\text{ISW}}(\hat{\mathbf{n}}) = - \int e^{-\tau(z)} (\dot{\Phi} + \dot{\Psi}) [\eta, \hat{\mathbf{n}}(\eta_0 - \eta)] d\eta, \quad (1)$$

where η is conformal time, the dots represent conformal-time derivatives, τ is the optical depth, and $e^{-\tau(z)}$ is the CMB photon visibility function. In the best-fit cosmological model, consisting of cold dark matter and a cosmological constant (Λ CDM), such anisotropies are generated at early times during the transition from radiation to matter domination and at late times, when dark energy begins to dominate, so these two contributions are known as the early and late ISW effects.

The early effect is typically generated shortly after recombination, so it is peaked around $l \sim 100$, and its contribution to the total CMB (which is small in the standard Λ CDM case) can be used to constrain the energy content of relativistic species, such as the number of neutrino species and their masses (Ichikawa et al. 2008), the presence of hot-dark-matter candidates such as massive neutrinos and axions (Hannestad et al. 2010), and interacting dark energy models (Väliiviita et al. 2010).

We focus here on the late effect as a probe of dark energy; in this case, the amplitude of the perturbations is also small compared with the primary CMB, and the fluctuations are generated on the largest scales, meaning that the effect is well described by linear theory. However, there is a small additional contribution from the non-linear growth in clusters, known as the Rees-Sciama effect (Rees & Sciama 1968). For a review, see Cooray (2002); in more recent work, Smith et al. (2009) provide a comparison with N -body simulations and perturbation theory, and Cai et al. (2010) give a comparison of linear and non-linear effects on the reconstructed ISW maps from ray tracing of CMB photons through N -body simulations. Finally, Schäfer et al. (2011) quantify the parameter estimation bias due to this non-linear effect and show that it is small compared to the statistical uncertainty imparted by cosmic variance.

2.1.1 Cross-correlations

Despite the small amplitude of the late ISW anisotropies, they can be used to constrain dark energy, as their presence can be detected by cross-correlating the observed CMB with the local matter density, which traces the gravitational potential (Crittenden & Turok 1996). The CMB temperature anisotropy $\Theta_{\text{ISW}}(\hat{\mathbf{n}})$ is given by Eq. (1), and the galaxy density contrast $\delta_g(\hat{\mathbf{n}})$ can be calculated as

$$\delta_g(\hat{\mathbf{n}}) = \int b_g(\hat{\mathbf{n}}, z) \varphi(z) \delta(\hat{\mathbf{n}}, z) dz, \quad (2)$$

where $\delta(\hat{\mathbf{n}}, z)$ is the dark matter density perturbation (linear theory suffices as described above), b_g the (linear) galactic bias and $\varphi(z)$ is the normalised visibility function of the chosen galaxy survey. This enables us to calculate the cross power spectrum,

$$C_l^{Tg, \text{ISW}} = \frac{2}{\pi} \int dk k^2 P(k) W_l^{T, \text{ISW}}(k) W_l^g(k), \quad (3)$$

where $P(k)$ is the matter power spectrum (linear theory suffices as well), and the source terms are, if we consider only the ISW temperature anisotropies,

$$\begin{aligned} W_l^{T, \text{ISW}}(k) &= - \int dz e^{-\tau(z)} \frac{d}{dz} [\tilde{\Phi}(k, z) + \tilde{\Psi}(k, z)] j_l[k\chi(z)] \\ W_l^g(k) &= \int dz \tilde{b}_g(k, z) \varphi(z) \tilde{\delta}(k, z) j_l[k\chi(z)], \end{aligned} \quad (4)$$

where the tilde denotes Fourier transformation and the j_l are the spherical Bessel functions. The auto-power spectra for the galaxies C_l^{gg} and the CMB (either the ISW part only $C_l^{TT, \text{ISW}}$ or the full observable spectrum $C_l^{TT, \text{tot}}$) can also be calculated by using the relevant source terms accordingly. In this work we calculate all the theoretical predictions implementing the above equations into a modified version of the CAMB integrator code (Lewis et al. 2000), without using the Limber approximation.

Notice that in the Λ CDM model and in most of its variants, as long as secondary Doppler effects due to reionisation can be neglected (Giannantonio & Crittenden 2007), the only significant source of large-angle CMB-density correlation is the ISW effect: we will therefore use the simpler notation C_l^{Tg} for the cross-correlations.

2.1.2 Signal-to-noise

The maximum signal-to-noise (S/N) ratio which is achievable for the ISW is limited by the amplitude of the primordial CMB perturbations. For an idealised full-sky full-depth survey, one can write (Crittenden & Turok 1996)

$$\left(\frac{S}{N}\right)^2 \leq \sum_l (2l+1) \cdot \frac{C_l^{TT, \text{ISW}}}{C_l^{TT, \text{tot}}}; \quad (5)$$

for the current WMAP7 Λ CDM model (Komatsu et al. 2011) described below in Section 3, this limit amounts to $S/N < 7.6$, as shown in Fig. 1. A more realistic estimation which takes into account the limitations of a galaxy survey, such as its redshift distribution, its shot noise due to finite surface density n_s (in sr^{-1}) and its sky coverage f_{sky} , is given by (see e.g. G08; Cabré et al. 2007):

$$\left(\frac{S}{N}\right)^2 \simeq f_{\text{sky}} \sum_l (2l+1) \cdot \frac{(C_l^{Tg})^2}{(C_l^{Tg})^2 + C_l^{TT, \text{tot}} (C_l^{gg} + 1/n_s)}. \quad (6)$$




Figure 1. Theoretical spectra and maximum signal-to-noise of the ISW detection for the current WMAP7 best-fit model (Komatsu et al. 2011). The top panel shows the angular temperature power spectrum of the ISW (green, solid) compared with the total CMB (red, dashed) and the data from WMAP7 (blue) (Larson et al. 2011). The bottom panel shows the maximum signal-to-noise (per mode and cumulative) achievable given the same model as given by Eq. (5) (brown, dot-dashed), as well as the improvement using CMB polarisation given by Eq. (8) (blue, solid).

Applying this expression to detections coming from current single galaxy catalogues typically yields only moderate significance ($2 \lesssim S/N \lesssim 3$).

Early attempts were made to measure the two-point correlation between COBE CMB data and tracers of the LSS (Boughn & Crittenden 2002), but these were limited by the noise and resolution of the COBE data. However, the ISW effect was soon detected using the WMAP data by many different groups exploiting a range of techniques; the first detection by Boughn & Crittenden (2004) used the X-ray background from the HEAO satellite and radio-galaxies from the NVSS survey, and was followed quickly by many others as we review below.

2.2 Single catalogue measurements

The first ISW analyses were focused on detection at any significance, measuring the two-point correlations of the WMAP CMB data and galaxy catalogues. This analysis can be performed equivalently in the real or harmonic spaces, using cross-correlation func-

tions or cross power spectra. While these approaches are formally equivalent, in practice some small differences can arise.

2.2.1 Cross-correlation functions

In real space, the observable quantity is the cross-correlation function (CCF) between CMB temperature and galaxies defined as

$$w^{Tg}(\vartheta) \equiv \langle \Theta(\hat{\mathbf{n}}_1) \delta_g(\hat{\mathbf{n}}_2) \rangle, \quad (7)$$

where the average is carried out over all the pairs of directions in the sky lying at the given angular separation $\vartheta = |\hat{\mathbf{n}}_1 - \hat{\mathbf{n}}_2|$. This approach has the advantage of being computationally straightforward, since the sky masks are defined in real space and are easily treated; the main drawback is the high level of covariance between the measured data points.

Following the release of the WMAP data, several groups reported positive detections using a wide range of LSS catalogues: the first measurement by Boughn & Crittenden (2004) used a combination of NVSS radio galaxies and X-ray background data, and the NVSS result was independently confirmed by Nolte et al. (2004). For optical surveys, indications were seen by Fosalba & Gaztañaga (2004) using the APM survey, and the evidence was improved by Fosalba et al. (2003) and Scranton et al. (2003) by using Luminous Red Galaxies (LRGs) from the SDSS. Cabré et al. (2006) detected the effect using the main SDSS galaxy sample, which while shallower, contains more galaxies than the LRG sample; Rassat et al. (2007) re-examined the relatively shallow 2MASS infrared survey, which was originally analysed with the power spectrum method by Afshordi et al. (2004). In Giannantonio et al. (2006), we reported the highest-redshift detection of the ISW with a catalogue of quasars from the SDSS, which was later reanalysed by Xia et al. (2009).

Most of these papers report a positive detection at low significance, typically between 2 and 3σ . One exception is the analysis using the 2MASS data where, though it favours the expected ISW signal, the evidence is very weak because the sample is too shallow for an appreciable ISW signal; it also is believed to have significant contamination from the Sunyaev-Zel'dovich (SZ) effect.

2.2.2 Cross-power spectra

We can also attempt to measure the angular power spectra of the cross-correlation directly. The main advantage of this approach is the relative decorrelation of the different modes, which makes the localisation of the signal on different scales more straightforward; the drawback is that the estimator of the correlation involves the inversion of a matrix whose dimensionality is the number of pixels over which the maps are projected (N_{pix}), which is computationally challenging, especially in realistic cases where the geometry of the survey mask is complex. For these reasons, approximate methods are generally used (for details, see e.g. Padmanabhan et al. 2003; Efstathiou 2004; Hirata et al. 2004; Ho et al. 2008; Schiavon et al. 2012), which still yield considerably lower correlation between modes when compared to cross-correlation measurements.

In this way, positive detections were reported by Afshordi et al. (2004) using the 2MASS catalogue, who simultaneously fit a template for the SZ effect; this was recently revisited by Francis & Peacock (2010b) who found weaker evidence more consistent with Rassat et al. (2007) and other analyses. Padmanabhan et al. (2005) applied the cross-spectrum technique to a SDSS LRG sample, and found consistent significance levels

to the earlier work, at about 2σ . More recently, Goto et al. (2012) measured the correlation between WMAP and the Wide-field Infrared Survey Explorer (WISE) survey, finding a high ISW signal at $> 3\sigma$. While the WISE volume is ~ 5 times larger than 2MASS, and thus the expected signal is higher, such a high correlation is $\sim 2.2\sigma$ higher than the Λ CDM expectations. Future analyses with the upcoming larger WISE data releases will help to clarify this issue.

Given accurate covariance matrices, cross-correlation and cross-spectra measurements (using the same data), should yield identical results, as, in total, both measurements contain the same information.

2.2.3 Other techniques

The ISW effect has also been seen using a method based on a wavelet decomposition by Vielva et al. (2006); McEwen et al. (2007, 2008); Pietrobon et al. (2006), who explored its dependence on different wavelet shapes and scales. While the significance level was sometimes reported to have been enhanced with this technique, the resulting constraints on cosmology were comparable with the previous two methods.

As the ISW is maximum on the largest scales, it is affected by the local variance, i.e. by the particular realisation of the matter distribution given the power spectrum; this may bias the results. For this reason, more advanced methods were developed to subtract the local variance, e.g. by Hernández-Monteagudo (2008); Frommert et al. (2008). In the latter work, a Wiener filter reconstruction of the LSS was used as a template instead of the theoretical cross-correlation function; it was estimated that this method may increase the S/N ratio by 7% in average.

It is also possible to reconstruct the ISW temperature maps. This was attempted by Barreiro et al. (2008) using NVSS data and a Wiener filter method, and by Granett et al. (2009) using SDSS data (see Section 5.1.6). Another optimised method has been later introduced by Dupé et al. (2011), based on the analysis of the temperature and density fields themselves rather than their spectra. A useful byproduct of this procedure is that a map of the ISW signal in the CMB is obtained. These authors also highlight the importance of separating the different statistical analyses, defining different procedures for testing the detection of a correlation in a model-independent way, measuring the confidence level based on a template, and comparing different models. This method was then validated with the 2MASS data, recovering a weak positive detection.

Another strategy to improve the signal-to-noise is to use the CMB polarisation information to reduce the primary CMB anisotropies, as proposed by Crittenden (2006); Frommert & Enßlin (2009); Liu et al. (2011). Depending on the details of the method, different authors estimate a level of improvement in the significance of the ISW detection in the range between 5 and 20%. In more detail, the correlation C_l^{TE} between CMB temperature and polarisation (E -modes) can be used to reduce the amount of primary anisotropies in the total temperature spectrum. Assuming idealised data, the maximum signal-to-noise of Eq. (5) is thus increased to

$$\left(\frac{S}{N}\right)^2 \leq \sum_l (2l+1) \cdot \frac{C_l^{TT, \text{ISW}}}{C_l^{TT, \text{tot}} - (C_l^{TE})^2 / C_l^{EE}}. \quad (8)$$

For the current WMAP7 Λ CDM model (Komatsu et al. 2011), we

find that this improves the upper limit to $S/N < 8.7$, as shown in Fig. 1.

Most of the approaches effectively measure a two-point statistic of the average correlations between CMB and LSS over the whole region covered by the surveys; however, with wavelets it is possible to try to localise sources of the ISW effect. This was more directly attempted by Granett et al. (2008), who identified massive superclusters and voids of galaxies in the SDSS LRG survey and their corresponding regions from WMAP were stacked to maximise the signal. A high significance detection was reported (at 4.4σ from this LRG catalogue alone), although this result was strongly dependent on the number of superclusters and voids used. See below Section 5.1.6 for a more detailed discussion.

2.3 Multiple catalogues measurements and their applications

The significance of the ISW detections can be increased by combining measurements obtained with multiple catalogues, to improve from a simple detection to parameter estimation and model comparison. Gaztañaga et al. (2006) made a first attempt at collecting all the existing measurements and used the resulting compilation to constrain cosmology; this was also extended by Corasaniti et al. (2005).

The difficulty with combining multiple measurements is achieving a reliable estimation of the covariance between them. It was proposed that a full tomographic analysis should be performed (Pogosian et al. 2005), including all the signals, and their covariances, as a function of redshift. This was finally achieved independently by Ho et al. (2008), using the harmonic space estimator, and by Giannantonio et al. (2008b) in real space, using five¹ and six galaxy catalogues respectively, summarising the state of the art in the field and upgrading the significance to 3.7σ and 4.5σ respectively. These results have been used to test a variety of dark energy and modified gravity models (Giannantonio et al. 2008a; Giannantonio et al. 2010; Lombriser et al. 2009; Lombriser et al. 2012; Lombriser 2011; Väliiviita & Giannantonio 2009; Serra et al. 2009; Daniel et al. 2009; Zhao et al. 2010; Bertacca et al. 2011); in the modified gravity case, the ISW provides a particularly useful constraint because it is sensitive to any non-trivial evolution of the gravitational potentials and the effective anisotropic stress.

2.4 Potential issues

Alongside these developments, some studies have questioned individual aspects of the ISW measurements, raising some doubts about the significance of its detection.

In Francis & Peacock (2010a,b), the 2MASS-CMB cross-correlation was re-analysed, and it was found that there is little evidence for an ISW detection from this catalogue alone, which is in agreement with most previous literature. But more worryingly, these authors also state that the ISW signal may remain undetected in 10% of cases (see Section 5.1.3 below).

In Hernández-Monteagudo (2010), the NVSS catalogue was re-considered, looking at its auto- and cross-correlation functions in both real and harmonic spaces. In both cases, some cross-correlation was seen, but the paper expressed concerns regarding a lower than expected signal on the largest scales and anomalous

¹ In the analysis by Ho et al. (2008) some of the catalogues were subdivided further into sub-samples.

large-scale structure in the NVSS map. We discuss these issues in Section 5.1.5.

Sawangwit et al. (2010) reanalysed some of the earlier ISW measurements and extended the analysis with three new LRG data sets, including a high-redshift sample developed using AAOmega spectra. The two catalogues at lower z were found to be in general agreement with a positive ISW signal, although at lower significance than seen elsewhere in the literature, while the high-redshift AAOmega sample showed no significant correlation. These authors also discussed the effect of possible systematics, suggesting that there is evidence of strong residual systematics from the study of data generated by rotating the real maps. We explore the rotation tests for our data and for the S10 data in Section 4.2; we then discuss the new data sets by S10 in Section 5.2.

Finally, López-Corredoira et al. (2010) reviewed some of the correlation analyses, finding levels of signal comparable to previous measurements, but significantly higher levels of uncertainty. We discuss this further in Section 5.1.4.

3 UPDATED DATA SET

We have updated our data compilation from G08 to include the latest available data for both the CMB and the LSS. All the data sets are pixellated in the Healpix scheme (Gorski et al. 2005) at a resolution of $N_{\text{side}} = 64$, corresponding to a pixel side of 0.9 degrees, as previously done in G08. We have checked that higher resolutions give consistent results.

In the following analysis, unless otherwise stated, we assume a fiducial flat Λ CDM cosmology, which we here update to the latest best-fit model from WMAP7+BAO+ H_0 , defined by energy densities for baryons $\omega_b = 0.0226$, cold dark matter $\omega_c = 0.1123$, sound horizon at the last-scattering surface $100 \vartheta_* = 1.0389$, optical depth $\tau = 0.087$, spectral index and amplitude of primordial scalar perturbations $n_s = 0.963$, $A_s = 3.195$, referred to a pivot scale $k_{\text{pivot}} = 0.002 \text{ h/Mpc}$ (Larson et al. 2011; Komatsu et al. 2011).

We summarise in Table 1 the most important properties of the data sets we use.

3.1 CMB data

The original data set by G08 was obtained by analysing the maps from the third year of WMAP, and it was checked upon the release of the WMAP five-year data that it yielded consistent results, as mentioned in Section IV.B of G08. Here we have updated the whole analysis by using the latest WMAP7 data (Jarosik et al. 2011), which should give a more stable foreground subtraction and noise reduction due to the increased integration time.

As for the choice of frequency, we use the internal linear combination (ILC) map and we also use the most aggressive galaxy mask associated with it, again on the basis that this should include the best foreground subtraction. We have checked that the signal does not change significantly between the different WMAP data releases, and that it is reasonably frequency-independent and close to the ILC result in the range of the WMAP bands Q, V, W. We discuss this further below.

3.2 Main SDSS galaxy data

The main galaxy distribution from the SDSS has been extended from Data Release Six (DR6) to the final imaging SDSS-III (DR8)

public data release (Aihara et al. 2011). These galaxies have been selected using the same criteria as in G08, i.e. starting from the photo- z primary galaxy sample (mode=1, type=3), which contains 208 million objects, and then imposing a cut in redshift of $0.1 < z < 0.9$ and a cut in flux of $18 < r < 21$, where r is the r -band model magnitude corrected for extinction. Also, only objects with photo- z uncertainty of $\sigma_z(z) < 0.5 z$ were considered. This leaves us with ~ 40 million galaxies, with a redshift distribution centred around $z \approx 0.3$. As the distribution of the photo- z 's can occasionally be inaccurate, for the calculation of the theoretical predictions we used a fit to a distribution function of the form introduced by Smail et al. (1994) and given by

$$\varphi(z) = \frac{1}{\Gamma\left(\frac{\alpha+1}{\beta}\right)} \beta \frac{z^\alpha}{z_0^{\alpha+1}} \exp\left[-\left(\frac{z}{z_0}\right)^\beta\right], \quad (9)$$

where the best-fit values of the parameters are $\alpha = 1.5$, $\beta = 2.3$, $z_0 = 0.34$; this fitted redshift distribution is similar to the DR6 result.

The mask was derived from a higher number density sample of the SDSS galaxies, selected with the weaker conditions $r < 22$ and $\Delta z < z$ (120 million objects), which was pixellated at a higher resolution ($N_{\text{side}} = 512$) and finding the number of filled high-resolution sub-pixels in each low-resolution pixel. Each low-resolution pixel was then assigned a weight f_i^g proportional to the fraction of high-resolution pixels which were filled.

An additional subtlety here is to avoid biasing the mask due to the high-resolution pixels which are on the edge of the survey themselves. We have found that the count-in-cells distribution of the 120 million galaxies in the high-resolution pixels is well approximated by a log-normal distribution of median $\mu = \ln(110)$, and is even better fit by the gravitational quasi-equilibrium distribution (GQED), described by Yang & Saslaw (2011). By comparing the best-fitting GQED distribution with the data, we found that the survey's edges introduce an enhanced tail in the distribution at low occupation number which leads to a bias in the mask. For this reason, we remove from the mask all high-resolution pixels with $n < 40$, since the best-fit GQED is nearly zero below this point. We found that our results are not overly sensitive to reasonable differences in the value chosen for this threshold.

We then mask the sky areas most affected by galactic extinction, dropping all pixels where the median extinction in the r band is $A_r > 0.18$. We have found that increasing this cut to the stricter level of $A_r > 0.16$ only changes the observed CCF by 5%. The unmasked survey area increased from 7,771 pixels (or 16% of the sky) for DR6 to 11,715 (or 24%) for DR8. The DR8 is the first data release to include a significant fraction of data from the Southern galactic hemisphere. We have checked that no significant difference appears when excluding data from the Southern hemisphere: the differences in the observed CCF are at the 10% level.

We estimated the galactic bias by fitting the Λ CDM prediction to the observed auto-correlation function (ACF), and found a value $b = 1.2$ assuming a scale- and redshift-independent bias.

3.3 Luminous Red Galaxies

We update the catalogue to include the latest data release of the MegaZ LRGs by Thomas et al. (2011b), which corresponds to the SDSS DR7, increasing our previous DR6 coverage by 10%. We apply the completeness cut in the de-reddened deVacouleurs i magnitude suggested by the authors ($i_{\text{dev}} - A_i < 19.8$) and we limit the star-galaxy separation parameter to $\delta_{sg} > 0.2$. We finally ap-

Catalogue	band	N after masking	f_{sky}	n_s [sr $^{-1}$]	\bar{z}	b
2MASS	IR	415,459	0.531	$6.23 \cdot 10^4$	0.086	1.3
SDSS gal DR8	Optical	30,582,800	0.253	$9.60 \cdot 10^6$	0.31	1.2
SDSS LRG DR7	Optical	918,731	0.181	$4.03 \cdot 10^5$	0.50	1.7
NVSS	Radio	1,021,362	0.474	$1.72 \cdot 10^5$	1.05	1.8
HEAO	X	N/A (flux)	0.275	N/A (flux)	0.90	1.0
SDSS QSO DR6	Optical	502,565	0.168	$2.39 \cdot 10^5$	1.51	2.6

Table 1. Summary of the properties of the LSS catalogues used. We report the number of objects after masking N , the sky fraction f_{sky} , the surface density of sources n_s , the median redshift of their distributions \bar{z} and the galactic bias b assumed constant needed to fit the auto-correlation function assuming the WMAP7 cosmology.

ply the reddening mask and discard pixels of median extinction at $A_r > 0.18$ as in the main galaxy case. As above, increasing this cut to the stricter level of $A_r > 0.16$ only changes the observed CCF by 5%. The redshift distribution in this case peaks around $z = 0.5$ and is smooth, and we use it directly as done in G08. The mask for the LRGs is that provided by Thomas et al. (2011b), with the addition of the aforementioned extinction mask. The bias of this catalogue is found to be $b = 1.7$, again by fitting to the measured ACF.

3.4 QSO data

For the quasars, no updated catalogue is yet publicly available, and here we use the same DR6 catalogue (Richards et al. 2009) as in G08, limiting ourselves to the cleaner subset of *UVX*-selected objects. To reduce the stellar contamination, which is more of an issue for quasars, we choose here a stricter extinction cut discarding pixels of $A_r > 0.14$. Increasing this cut to the stricter level of $A_r > 0.12$ changes the observed CCF by less than 10%.

In this paper we used Eq. (9) to determine a fit to the QSO redshift distribution, for the same reasons described in Section 3.2, and obtained best-fit parameters of $\alpha = 2.0$, $\beta = 1.5$, $z_0 = 1.06$, corresponding to a median $\bar{z} = 1.5$. This is changed from G08 where the visibility function was simply binned, and not smoothed, and so included irregular steps between adjacent bins.

The linear bias parameter found from the QSO ACF is $b = 2.6$, 10% higher than that reported in G08. The amount of stellar contamination, as seen by comparing the large-scale power in the ACF with the ACF of a catalogue of stars from the SDSS, is consistent with a fraction $\kappa = 2\%$, which is expected in these data (Richards et al. 2009).

3.5 Other data

For the other surveys (2MASS, HEAO, and NVSS), we continue using the same maps as in the previous analysis of G08. However, we have improved the analysis in the following ways:

For the low redshifts probed by 2MASS, non-linear effects are large enough to significantly affect the zero-lag bin of the ACF; this was seen by comparing the linear power spectrum with the result obtained by Smith et al. (2007) at the scale of one degree. We have therefore dropped this bin which is significantly higher than the linear theory prediction, so that the linear bias has decreased to $b = 1.3$ in this case.

For NVSS, a well-known issue is the uncertainty in the redshift distribution of the sources. Significant changes to the redshift

distribution can affect the cross-correlation template, and so impact the detection significance, and will also change the predicted Λ CDM cosmology amplitudes. Here we have compared the correlation functions obtained using the distribution function based on the models of Dunlop & Peacock (1990) (used in G08), to the distribution from De Zotti et al. (2010), who fitted the template based on radio galaxies with measured redshifts, and the distribution function introduced by Ho et al. (2008)², who simultaneously fit the cross-correlation functions between NVSS and other galaxy catalogues. As seen in Fig. 2, the theoretical predictions for these models are compatible within the expected measurement error bars. We find that the resulting significance of the NVSS ISW detection changes at most by 10% when using each of the three models for the source distribution. In the following, we use the model by De Zotti et al. (2010), as it is based on a subsample of galaxies of known redshifts. Further, we have included a better modelling of the shot noise in the ACF, due to the fact the maps were originally pixellated at a lower resolution. This primarily affects the ACF, and so the measurement of the bias; we now find $b = 1.8$, increased from $b = 1.5$ previously assumed in G08.

For the HEAO catalogue, we have also included the same pixellation correction in the shot noise modelling, but as the instrumental beam is much larger ($\theta_{\text{FWHM}} = 3.04$ degrees), the resulting bias parameter remains $b = 1.0$.

3.6 Method

We pixellate all these data on the sphere as described above, and measure the two-point functions between them and the CMB, using the simple estimator

$$\hat{w}^{Tg}(\theta) = \frac{1}{N_\theta} \sum_{i,j} \frac{(n_i - \bar{n})}{\bar{n}} (T_j - \bar{T}) f_i^g f_j^T, \quad (10)$$

where n_i and T_i are the number of galaxies and the CMB temperature in a pixel of masked weights f_i^g, f_j^T respectively, \bar{n}, \bar{T} are the average number of galaxies per pixel and average temperature, and $N_\theta = \sum_{i,j} f_i^g f_j^T$ is the weighted number of pairs at a given separation. Note that in our approach the CMB weights f_j^T are simply taken to be either 0 or 1, depending on whether the pixel is masked or unmasked. When $f_i^g < 1$, which occurs mostly at the edges of the surveyed area, the number of galaxies in a pixel n_i needs to be rescaled from the observed number n_i^{obs} as $n_i = n_i^{\text{obs}} / f_i^g$, in order to

² Note the typo in Eq. (33) of Ho et al. (2008), where the argument of the Gamma function should be $(\alpha + 1)$ to ensure the stated normalisation of $\int f(z) dz = b_{\text{eff}}$.




Figure 2. The redshift distribution of NVSS and its effect on the ISW correlation. In the top panel we show three normalised selection functions which have been used widely in the literature, and in the bottom panel the resulting theoretical cross-correlation functions, compared with our data, for the fiducial Λ CDM model. The bias is constant, and set to $b = 1.8$ for the first two models as this gives good agreement with the ACF. It is $b = 1.98$ for the last model, as fitted by Ho et al. (2008). We can see that the differences are small compared with the measurement error bars. We use the model by De Zotti et al. (2010) in the main analysis.

keep the same mean density \bar{n} over the whole map. As in G08, we use 13 angular bins linearly spaced between 0 and 12 degrees.

We estimate the full covariance matrix C using the ‘MC2’ Monte Carlo method described in G08; specifically, based on the fiducial flat Λ CDM cosmology, we generate Gaussian random maps of the CMB and of all the galaxy catalogues, using their known redshift distributions, number densities, and including all

the expected correlations between the catalogues. We also add the expected level of Poisson noise based on the surface densities of each catalogue on top of all realisation of the Gaussian maps. For each of 5000 realisations of these maps, we measure the correlation functions, and calculate the covariance of them. (See the Appendix of G08 for more details.) We have confirmed that 5000 realisations are enough for convergence of the signal-to-noise. As the cross-correlations are in agreement with the fiducial Λ CDM cosmology used in the mocks, we expect this modelling of the covariance should be reasonably accurate. See below Section 5.1.1 for more details and a comparison with the analytic covariance. Most catalogues are significantly covariant; the samples which are less covariant with the others are the LRGs and the QSOs, because of their unique redshift coverage, which is more peaked in the former case, and deeper in the latter.

We fit the amplitudes assuming the cross-correlation functions are Gaussianly distributed; this is approximate, as even if the maps themselves are Gaussian, as assumed in our Monte Carlos, two point statistics of those maps will not be Gaussian. However, this appears to be a reasonable approximation for correlation functions (and particularly cross-correlations), where each bin represents an average of many products of pixels and the central limit theorem should apply. This is confirmed in our Monte Carlos, where we find the skewness in the covariance to be relatively small, with dimensionless skewness measure of 0.1 – 0.2. However, it may be worth further investigating any residual bias which this level of non-Gaussianity might cause in future work. Non-Gaussianity is likely to be a more significant concern for power spectrum estimators, particularly for auto-spectrum measurements which must be positive-definite.

3.7 Results and public release

The results of the new cross-correlation analysis are shown in Fig. 3, and are in general agreement with G08 given the measurement errors. We can see that all the measurements lie close to the Λ CDM prediction; however the LRGs do show an excess signal at the $> 1\sigma$ level. See Section 4 for a discussion of possible systematic effects.

The only CCFs for which we see a non-negligible change in Fig. 3 compared to the earlier analysis by G08 are the LRGs and the NVSS, where the signal has somewhat increased. This appears to be primarily due to changes in the WMAP data rather than in the LSS surveys; in Fig. 4 we show the CCFs resulting when the current LSS maps are correlated with different WMAP data releases. With the exception of the LRGs, the changes tend to bring the data into better agreement with the Λ CDM theory. We have found that similar changes appear if the single-frequency, cleaner maps (V and W) are used instead of the ILC. Further, we found that a significant part of these changes is induced by the change in the WMAP mask between the different releases rather than the change in the data themselves, suggesting that the differences may be due to a better foreground cleaning.

To provide a global estimate of the combined significance, we use a theory template $\bar{w}^{ISW}(\theta_i) = Ag(\theta_i)$, where $g(\theta_i) \equiv g_i$ is the theoretical prediction of the WMAP7 best fit model and A is the fit amplitude for each catalogue; further details can be found in G08.




Figure 3. Updated results of the cross-correlation of all the data sets with the WMAP7 ILC map. Most data (dark green, solid) are in good agreement with the theoretical predictions for a Λ CDM model (red, short-dashed), with the only exception of the LRGs which show an excess at $>1\text{-}\sigma$ level. The highly correlated error bars are from 5000 Monte Carlo mocks and are $1\text{-}\sigma$, except for 2MASS where they are $0.5\text{-}\sigma$ to improve readability of the plot. Further, the first five data points for 2MASS have been excluded due to potential contamination by the SZ effect. The light-green, long-dashed lines show the previously published data by G08, and the blue, dot-dashed lines are the best amplitude fits.




Figure 4. Dependence of the results on the different WMAP data releases. Most results vary little compared with the size of the error bars; the largest changes can be seen in NVSS, bringing the results closer to the Λ CDM expectations. The error bars on 2MASS are again $0.5\text{-}\sigma$.

By analytically maximising the likelihood, we obtain that the best value A and its variance for each catalogue are:

$$A = \frac{\sum_{i,j=1}^p C_{ij}^{-1} g_i \hat{w}_j^T}{\sum_{i,j=1}^p C_{ij}^{-1} g_i g_j}, \quad \sigma_A^2 = \left[\sum_{i,j=1}^p C_{ij}^{-1} g_i g_j \right]^{-1}, \quad (11)$$

where \hat{w}_i^T are the observed CCF for each survey (sampled in $p = 13$ angular bins) and C_{ij} is the measured covariance matrix of dimension p described above. To obtain an unbiased estimator of the inverse covariance C_{ij}^{-1} , we correct the result obtained by inverting C_{ij} by a factor $\alpha = (N - p - 2)/(N - 1)$ (Hartlap et al. 2007); however in our case (for $N = 5000$ realisations) this correction is negligibly small. This method can be immediately generalised to the full case, in which we fit a single amplitude to a template which includes the six CCFs. In this case the total number of angular bins, and thus the dimension of the covariance matrix, becomes $p = 6 \times 13 = 78$.

The results with this method and the new data are given in Table 2, where we can see that if we identify the signal-to-noise ratio as $S/N = A/\sigma_A$, then the total significance of a detection is now at the 4.4σ level when a single amplitude is used for all six catalogues.

Catalogue	$A \pm \sigma_A$	S/N	expected S/N
2MASS cut	1.40 ± 2.09	0.7	0.5
SDSS gal DR8	1.24 ± 0.57	2.2	1.6
SDSS LRG DR7	2.10 ± 0.84	2.5	1.2
NVSS	1.21 ± 0.43	2.8	2.6
HEAO	1.37 ± 0.56	2.4	2.0
SDSS QSO DR6	1.43 ± 0.62	2.3	1.7
TOTAL	1.38 ± 0.32	4.4 ± 0.4	$\approx 3.1, < 7.6$

Table 2. Results from the updated data set compared with the expected signal-to-noise calculated using Eq. (6) for each catalogue. The first five data points for 2MASS have been excluded. For the total expected S/N , we show both the value estimated from the MC mocks (see below), and using the upper limit of Eq. (5). We estimate a $\sim 0.4\sigma$ systematic error on the total signal-to-noise ratio, due to possible different masks and other choices entering into its determination.

It is worth noticing that our significance estimation is however based on a fiducial model which includes not only the WMAP7 best-fit parameters, but also the assumed redshift distributions and simple bias model of the surveys. While this is reasonable to give an initial estimate of the significance of the detection, a full cosmological analysis should ideally take into account the uncertainties in these quantities, e.g. with the help of additional nuisance parameters. The assumption of constant biases is especially uncertain, in particular for very deep catalogues like HEAO and NVSS (see e.g. Schäfer et al. 2009), and this issue should be addressed in a full cosmological analysis allowing for a more realistic bias evolution. The different assumptions for the biases and for the redshift distributions may for example explain the difference between our results and those by Ho et al. (2008), where a higher excess signal was found, at the 2σ level above the Λ CDM predictions.

In Table 2 we also compare the results with the expected signal-to-noise calculated using Eq. (6) for each catalogue, and using the upper limit of Eq. (5) for the total. We can see that the measured results are higher than the expectations in most cases. Given this discrepancy exists between expectations and observations, we next proceed to quantify its significance by studying the distribution of the ISW signal-to-noise obtained using our 5000 mock maps of the galaxy surveys. We show these distributions in Fig. 5: here we can see that they are broad, and the position of the observed S/N is well within the expected scatter. In more detail, we fit a Gaussian to the distribution of the mock total signal-to-noise, where we find that the mean (i.e. the expectation for the total S/N from Λ CDM) is 3.05, and the r.m.s. is 1 by construction. This places our observed result 1.35σ away from the mean.

A further interesting point to be learnt from Fig. 5 is the comparison between theoretical S/N with (green solid lines) and without (green dashed lines) shot noise. We can see that the effect of shot noise is particularly large for the quasars, due to their limited number density. From this we conclude that future measurements of extended quasar catalogues have the potential to significantly improve the existing results, due to the large redshift coverage of these sources.

Given the number of different assumptions in the method of the analysis, we roughly estimate that a systematic uncertainty of ≈ 0.4 needs to be included on the final figure of the signal-to-noise ratio. For example, using other reasonable redshift distributions for the catalogues typically results in changes of the signal-to-noise ratio of the order $0.2 - 0.4$. Similar differences are obtained when

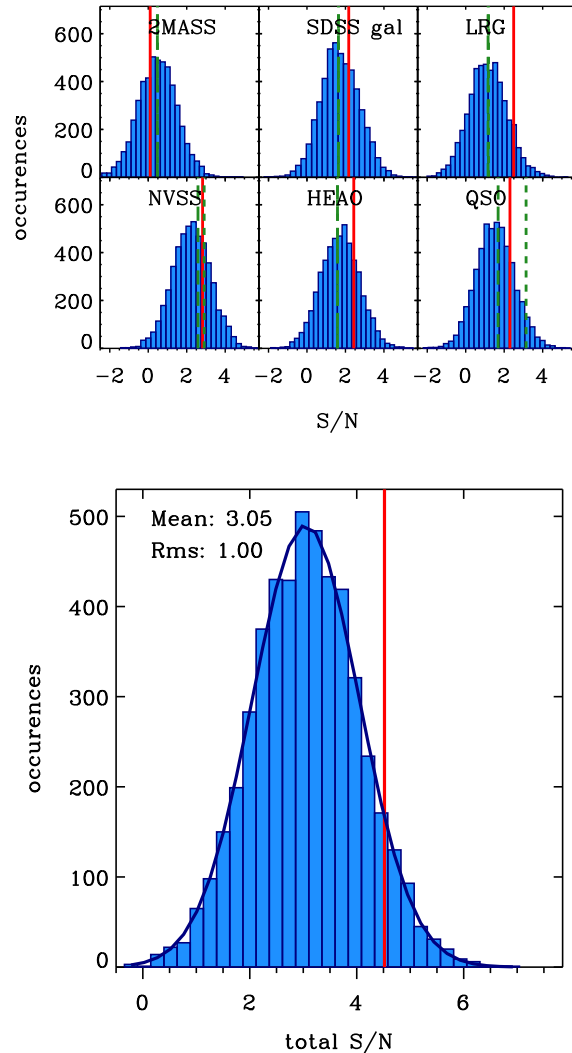


Figure 5. Distribution of the signal-to-noise ratio for our 5000 MC realizations (blue histograms), compared with the observations (red solid lines) and the theoretical expectations (green). The different green lines refer to: no shot noise (short-dashed), and shot noise included (long-dashed). The top panel shows each catalogue separately, the bottom panel is the full combination, for which we also show the best-fit Gaussian distribution and its parameters.

changing the thresholds in the extinction masks, or excluding parts of the data (such as the Southern hemisphere for the new SDSS DR8 galaxies). Another change which is typically at the same level is produced if we decide to completely discard the pixels near the edge of the survey, for which the mask weighting is $f_i^s < 1$, instead of correcting them with the appropriate weights. Furthermore any extra large-scale power in the auto-correlation functions, which could arise e.g. due to low-redshift contamination or other systematics, would increase the variance of the cross-correlations. We discuss this below in Section 5.3.2, showing that its effect is limited and in agreement with our systematic estimation of 0.4σ .

We have also checked the effect of removing any one catalogue from the analysis, finding in this case that the total significance can not be lowered below 3.9σ , which is the result obtained when ignoring the NVSS data.

We publicly release the maps, masks, and results discussed herein, which can be downloaded from the internet.³ Here one can find for each of the six catalogues the following data: a file with the Healpix galaxy map in FITS format, and its companion mask in the same format; a table with the redshift distribution, and a table with the results of the CCFs. Finally, the full covariance matrix based on Monte Carlo maps is also provided.

4 SYSTEMATIC UNCERTAINTIES

Since the ISW signal is expected to be weak, possible systematic uncertainties are a serious issue. Here we discuss some of the new tests we have explored to constrain this contamination, and further discussion can be found in G08 and earlier work.

4.1 Foreground contamination

While individual instrumental systematics are not expected to be correlated between surveys, the cross-correlation could be contaminated by extra-galactic sources in the microwave frequencies, such as synchrotron emission or the Sunyaev-Zel'dovich effect. Our galaxy also emits in the microwave (dust, synchrotron and free-free), so it is important to ensure that galactic structure does not creep into the large-scale structure measured in surveys; otherwise spurious correlations with the CMB foregrounds will bias the cosmological interpretation of the measurements. Another possible source of cross-correlation is the secondary Doppler effect which may be added to the CMB at the time or reionisation (Giannantonio & Crittenden 2007); however, at the redshifts of currently available data, this is expected to be completely subdominant.

An important way to keep foreground contamination under control is to check for dependence of the cross-correlation functions on the CMB frequency, as we expect the ISW signal to be monochromatic while the extra-galactic and galactic foregrounds sources such as SZ, synchrotron and dust are expected to have a strong frequency dependence. We have performed this test for the new data, and the results can be seen in Fig. 6; the cross-correlation functions are quite stable relative to the measurement error bars for the less contaminated WMAP maps (ILC, W, V, and largely Q). The most dependent on frequency is 2MASS, where Afshordi et al. (2004) suggested there was evidence for the SZ effect; the most stable is HEAO, which is perhaps understandable as the hard X-ray background should be least affected by galactic contamination.

Perhaps the most worrying systematic problem which can affect the ISW measurement is the leakage of information from the structure of our Milky Way into the galaxy catalogues. This can potentially jeopardise the results, as it correlates with the CMB via residual dust extinction corrections. This problem can be minimised by masking sky areas closest to the galactic plane, and those areas most affected by reddening, which we have done for all catalogues.

We have also checked for the effect of removing from the maps a band centred on either the ecliptic or galactic planes, using cuts of different widths (10, 20 and 30 degrees). We have found no significant differences. Other possible systematics include the effects of point sources, regions of poor seeing and sky brightness;




Figure 6. Frequency dependence of the cross-correlations, for WMAP7: most results are stable compared with the size of the error bars. The error bars on 2MASS are again 0.5σ .

as discussed in G08 they are less severe than dust extinction for our data.

For the DR7 MegaZ LRG data set there is excess power in the correlation with the CMB, at a similar level as found with DR6; excess large-scale power has also been seen in the MegaZ LRG auto-correlation function as discussed by Thomas et al. (2011a). While these excesses may be cosmological, they could also be evidence of remaining systematics, such as a higher than expected stellar contamination. Recently, Ross et al. (2011b) improved the methods for selecting and interpreting an LRG catalogue, using a large spectroscopic training set from SDSS-III BOSS (Eisenstein et al. 2011). These authors found that the most serious systematic problem with present photometric redshift catalogues from SDSS imaging is the failure to detect faint galaxies around foreground stars (even relatively faint stars to $r \approx 20$). Ross et al. (2011b) corrected this issue, but it does illustrate the need to be diligent about stellar contamination, especially when cross-correlating with other datasets which may also include some contamination with galactic sources. In the future, catalogues with this higher level of systematics control should be used for the measurement of the ISW. The ISW signal seen in the LRG cross-correlation is higher than the Λ CDM expectation; if future LRG data will be more in accord with Λ CDM, the best-fit amplitude may well decline, but with the uncertainty also shrinking, the signal-to-noise ratio could remain at a similar level.

4.2 The rotation test

Although the presence of systematic effects has been studied rather carefully by several authors, it is possible some unaccounted uncertainty could remain in the data, thus compromising the measurements, and it is therefore worthwhile to look for new ways to check for such problems.

One such test, based on arbitrary rotations of the sky maps, was recently proposed by Sawangwit et al. (2010). In this test, cross-correlation functions (CCF) are generated by cross-correlating the true maps, but after one of the maps has been *rotated*

³ www.usm.lmu.de/~tommaso/isw2012.php




Figure 7. The rotation test for our data, zero-lag case. In the top panel, for each galaxy catalogue, we show the evolution of the CCF at 0 degrees as a function of the arbitrary rotation angle $\Delta\varphi$, describing rotations around the galactic plane. The coloured lines show the observed results; the green shaded bands indicate 1 and 2 σ regions calculated by generating 500 mock Monte Carlo maps of the data. The black dashed lines are the averages seen in the same mock data, assuming a Λ CDM concordance model. The bottom panel shows the same test for rotations in different coordinate systems: galactic, equatorial, ecliptic and supergalactic.

by some arbitrary angle $\Delta\varphi$ relative to the other. In Sawangwit et al. (2010) rotations around the Galactic axis were chosen particularly to try to identify galactic contamination, but any axis choice is possible. For a sufficiently large rotation $\Delta\varphi$, one expects on average that there will be no correlation, though any particular measurement will have scatter reflecting the intrinsic variance in the measurement. (Any rotation leaves two poles fixed, implying a small amount of residual correlation, but in practice this is negligible.)

The critical question is, how do we evaluate the significance of the rotated correlation functions? The most obvious comparison is to the variance of the intrinsic scatter, inferred from our Monte Carlo simulations. Are the rotated correlation measurements consistent with this intrinsic scatter or not?

4.2.1 Assessing the significance

One difficulty in making this comparison is that, for a given rotation axis, there are a limited number of rotations which are independent of one another due to the fact that the ISW signal is on relatively large angles. We use 500 Monte Carlo simulations to evaluate this; we create sets of CMB and galaxy maps with the expected signal for the Λ CDM model, and rotate them as we do the real data. In the top panel of Fig. 7 we show the average ‘zero-lag’ cross-correlation as a function of rotation angle, with the 1- σ and 2- σ regions shaded in green. We see that this is significant typically out to 30 degrees, implying that this is the minimum rotation required to provide an independent sample.

For any given rotation axis, this leaves only 11 independent samples of the cross-correlation measure. If one or more of these are significantly higher than the r.m.s. amplitude, then this could be evidence for the systematic contamination. In the top panel of Fig. 7 we plot the ‘zero-lag’ cross-correlation for each of our catalogues. This shows that no significant outliers exist, and that the signal is generally highest when there is no rotation, apart from 2MASS where the zero-lag signal could be cancelled by the SZ effect. This is quantified in Table 3, where we count the number of rotations exceeding thresholds of 1-, 2 σ (expected 32% and 5%); if anything, the number of high correlations seen in the rotated maps is lower than expected, though this discrepancy is not significant. The observed number above the threshold should be binomially distributed, which is very broad given the limited number of observations.

In S10, they looked instead at the number of rotations where the ‘zero-lag’ cross-correlation exceeded the true ‘zero-lag’ cross-correlation for that particular survey; that is, our zero degree value (denoted w_{0i}). However, this is just one choice, with the drawback of being different for each data set i . Nonetheless, we also show this in Table 3. Again, for rotations about the galactic axis, we find none exceeding the true value for our data sets.

We can increase our statistics and explore for the possibility of systematics associated with other reference frames by considering rotations about other axes. We have repeated the test with rotations using: galactic, ecliptic, equatorial and supergalactic axes. For the discussion below we ignore chance alignments of the various rotations, but it should be kept in mind that all these points may not be fully independent. The results are shown in the bottom panel of Fig. 7. *All of the rotations are consistent with the expected scatter, showing no indication for a preferred rotation axis.* This seems to support the hypothesis that the rotations are merely providing a measure of the intrinsic scatter and are not suggestive of a contaminant associated with a single axis. The statistics are summarised in Table 3, where we can see that only 31% and 3% of the points lie above the 1-, 2- σ thresholds respectively. Excluding 2MASS, only 3% exceed the true value measured in the data; unfortunately, this statistic cannot be used directly to estimate a total significance because it will be dominated by maps with the least significant amplitude.

For a visual impression of this, we plot our results in Fig. 8, where the observed number of detections in excess of each threshold is compared with the 68% confidence regions drawn from the continuous generalisation of a binomial probability distribution. Here we can appreciate that the distribution of the points above the 1- and 2- σ thresholds is fully consistent with the expected scatter and the number of points above the true measure levels seems consistent with an average level of significance $\sim 2\sigma$ for each catalogue.

Zero-lag data only						
Catalogue	$> 1\sigma \neq 0$	– multi-axes	$> 2\sigma \neq 0$	– multi-axes	$> w_{0i}$	– multi-axes
2MASS	0/11 (0%)	12/44 (27%)	0/11 (0%)	2/44 (5%)	—	—
SDSS gal	4/11 (36%)	16/44 (36%)	0/11 (0%)	1/44 (2%)	0/11 (0%)	3/44 (7%)
SDSS LRG	5/11 (45%)	16/44 (36%)	0/11 (0%)	3/44 (7%)	0/11 (0%)	2/44 (5%)
NVSS	3/11 (27%)	11/44 (25%)	0/11 (0%)	0/44 (0%)	0/11 (0%)	0/44 (0%)
HEAO	1/11 (9%)	10/44 (23%)	0/11 (0%)	0/44 (0%)	0/11 (0%)	0/44 (0%)
SDSS QSO	2/11 (18%)	17/44 (39%)	1/11 (9%)	3/44 (7%)	0/11 (0%)	1/44 (2%)
Total	15/66 (23%)	82/264 (31%)	1/66 (2%)	9/264 (3%)	0/55 (0%)	6/220 (3%)
Expected	21/66 (32%)	84/264 (32%)	3/66 (5%)	13/264 (5%)	—	—

Table 3. Results of the rotation test with our data for the zero-lag CCF. The scatter is in reasonably good agreement with the expectations, and is smaller than in S10 data (see Table 5). In the rotations around the galactic axis there are fewer outliers than expected, but the numbers become closer to the expectations when using a larger number of rotation axes.

Template amplitude fitting						
Catalogue	$> 1\sigma \neq 0$	– multi-axes	$> 2\sigma \neq 0$	– multi-axes	$> A_i$	– multi-axes
2MASS	0/11 (0%)	12/44 (27%)	0/11 (0%)	3/44 (7%)	—	—
SDSS gal	4/11 (36%)	17/44 (39%)	0/11 (0%)	3/44 (7%)	0/11 (0%)	2/44 (5%)
SDSS LRG	6/11 (55%)	17/44 (39%)	1/11 (9%)	2/44 (5%)	0/11 (0%)	0/44 (0%)
NVSS	3/11 (27%)	10/44 (23%)	0/11 (0%)	2/44 (5%)	0/11 (0%)	0/44 (0%)
HEAO	4/11 (36%)	15/44 (34%)	0/11 (0%)	0/44 (0%)	0/11 (0%)	0/44 (0%)
SDSS QSO	3/11 (27%)	18/44 (41%)	1/11 (9%)	3/44 (7%)	1/11 (9%)	3/44 (7%)
Total	20/66 (30%)	89/264 (34%)	2/66 (3%)	13/264 (5%)	1/55 (2%)	5/220 (2%)
Expected	21/66 (32%)	84/264 (32%)	3/66 (5%)	13/264 (5%)	—	—

Table 4. As above, for the template amplitude fitting. The scatter is here in even better agreement with the expectations.

Thus we find no evidence for systematic contamination from the rotation tests and that the rotated cross-correlation measurements are consistent with those expected from the covariance matrix derived from Monte Carlo simulations. Indeed, the rotations are an alternative means of deriving the covariance and were used in early studies as a means of calculating it. Unfortunately, the number of possible independent rotations is limited, meaning it is difficult to derive a stable covariance from rotations alone.

4.2.2 Including full angular information

For the discussion above, we have focused on the zero-lag cross-correlation following S10. However, this single bin contains only part of the ISW signal, and thus the significance derived from it is not as high as one gets when including the full cross-correlation function. For this reason, we repeat our analysis fitting to templates based on the predicted correlations expected in the WMAP7 best-fit cosmology. We use the covariance matrices obtained with the expected CCF; we plot the best-fit amplitudes of the rotated maps as a function of the rotation angle in the first panel of Fig. 9. The scatter of the A 's from the rotated maps is in even better agreement with expectations, as can be seen in Table 4. Here, we see that only 30% and 3% of the points lie above the thresholds of 1σ and 2σ . 2% exceed the level of the unrotated best-fit amplitude, denoted as A_i .

We have also performed the multi-axes test in the amplitude fitting case, as shown in the second panel of Fig. 9. We can see in

the figure, and in the results summarised in Table 4, that the results remain consistent over an increased population.

The significance of these results can be better appreciated by considering the uncertainties on the number counts, shown in the lower panel of Fig. 8. Here we can see once again that not only are the 1- and 2- σ thresholds consistent with the distribution of the outlying data, but also the populations above the A_i thresholds are consistent with a level of significance $> 2\sigma$ for each catalogue (excluding the 2MASS data, where the detection has low significance).

4.2.3 Application to Sawangwit et al. (2010) data

We have shown above that for our data, the results of the rotation test are in agreement with the Monte Carlo estimations of the variance; let us now discuss the application to the S10 data themselves. It was claimed by Sawangwit et al. (2010) that the result of this test for multiple data sets (SDSS gal, NVSS, SDSS LRGs and AAOMEGA LRGs) was in contradiction with the claimed detection of the ISW, pointing instead towards strong unknown residual systematics.

By looking at Figure 14 in S10, we compile the statistics shown in Table 5. Here we can see that in total 39% of the random points obtained by arbitrary rotations are scattered at $> 1\sigma$ away from zero, while randomly we would expect this fraction to be 32%. If we instead choose the 2- σ threshold, we find a total of 6/56 points above this threshold, corresponding to 11% compared to the expectation of only 5%. Alternatively following S10, we can

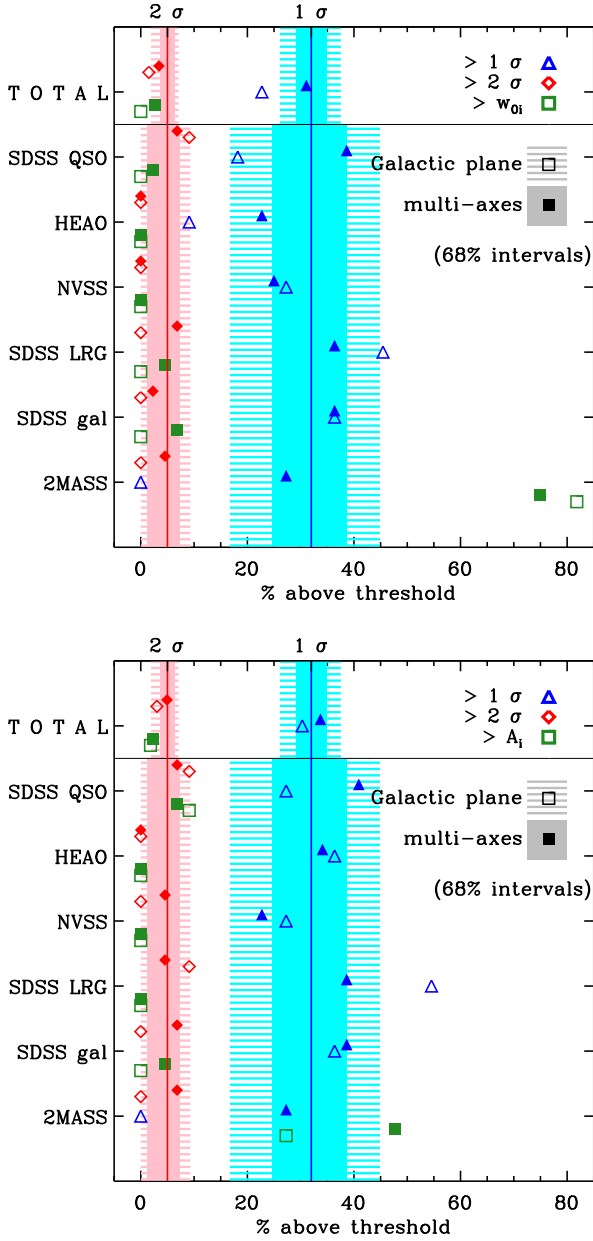


Figure 8. Significance of the rotation test for our combined data, using WMAP7, zero-lag only (above) and full template amplitude fitting (below). The data points correspond to the percentage of realisations observed above any given threshold: 1σ (blue triangles), 2σ (red diamonds), and each catalogue’s unrotated CCF w_{0i} or amplitude A_i (green squares). These should be compared to the expected 1- and 2- σ fraction (vertical lines); the error bands are 68% confidence intervals assuming a binomial probability distribution for the counts. The empty points and the dashed bands refer to rotations around the Galactic plane only, while the filled points and solid bands include data from all four rotation axes.

choose the threshold to be the unrotated signal w_{0i} seen in the data; in this case we find that 13/48 points (27%) are above this mark, if we discard the points from AAOMega for which S10 find no correlation.

By including the expected errors on the counts of the outlying




Figure 9. The rotation test for our data, best amplitude case. In this case we show in the top panel the best fit amplitude, using all the angular information, and assuming a templated based on the best fit WMAP7 Λ CDM cosmology; the bottom panel shows the same using multiple rotation axes: Galactic, Equatorial, Ecliptic and Supergalactic.

points, we can see the uncertainties corresponding to the data by S10 in Fig. 10. Here we plot, for each data set and for the total, the number of bins above each threshold of 1σ , 2σ , and the ‘real’ value of the CCF at no rotation w_{0i} (again, excluding the AAOMega data for this last case since here w_{0i} is not defined). We can see here a higher scatter than in our data, but this not unexpected given the lower number of realisation. By comparing the data with the 68% and 95% confidence intervals from the binomial distribution, we observe substantial agreement.

With respect to the w_{0i} threshold, Fig. 10 indicates that NVSS and, to a lesser extent the SDSS 20-21, 19-20, and 2SLAQ LRGs are consistent with a significance level $> 2\sigma$, while the other catalogues seem to have a lower significance level. Also, it should be mentioned that the error bars were calculated using a jack-knife method, which can often depend on the details of how the jack-knife was performed. In earlier studies, jack-knife error estimates have been seen to be somewhat smaller than those inferred from Monte Carlo methods (Cabr   et al. 2007). This could explain why somewhat more of the random rotations appear significant in the S10 data compared to our analysis. Overall however, when consid-




Figure 10. Significance of the rotation test for the data of S10, for each individual catalogue and (on top) for the total. The data points correspond to the percentage of realisations observed above any given threshold: 1σ (blue triangles), 2σ (red diamonds), and each catalogue’s unrotated CCF w_{0i} (green squares). The expected 1- and 2σ thresholds are overplotted with vertical lines. The error bands in this case are 68% (solid) and 95% (dashed) confidence intervals assuming a binomial probability distribution.

Catalogue	$> 1\sigma \neq 0$		$> 2\sigma \neq 0$		$> w_{0i}$	
SDSS LRGs	1/8	13%	0/8	0%	3/8	38%
2SLAQ LRGs	2/8	25%	0/8	0%	2/8	25%
AAOmega LRGs	4/8	50%	1/8	13%	—	—
SDSS gal $18 < r < 19$	5/8	63%	1/8	13%	3/8	38%
SDSS gal $19 < r < 20$	4/8	50%	2/8	25%	2/8	25%
SDSS gal $20 < r < 21$	3/8	38%	2/8	25%	2/8	25%
NVSS	3/8	38%	0/8	0%	1/8	13%
Total	22/56	39%	6/56	11%	13/48	27%
Expected	18/56	32%	3/56	5%	—	—

Table 5. Results of the rotation test in Sawangwit et al. (2010). By definition, we would expect 32% (5%) of the points to be further than 1σ (2σ) away from zero.

ered with the results based on our maps, we find that *the rotation test provides no significant evidence for residual systematics in the data.*

5 DISCUSSION

A handful of other papers have appeared with results that seem to be in conflict with the more common approaches described above. The differences are centred on three main issues: large differences in the statistical significance of the results (with papers suggesting both higher and lower significance), questions related to new data sets and questions about systematic contaminations. We discuss each of these in turn.

5.1 Statistical significance

5.1.1 Theoretical covariance

The significance of the measurements is estimated given the full covariance between the data. This may be obtained using different methods: analytically (“theory” covariance), from the data themselves (jack-knife or bootstrapping), or using Monte Carlos, which is our method of choice. It was shown already by Cabré et al. (2007) that all such methods should give consistent results (with fluctuations of order 10% between them), and G08 presented a comparison between MC and jack-knife (JK) covariances showing that the results were comparable, although the JK method was deemed less stable. Here we compare our baseline MC covariance with the theory (TH) covariance calculated analytically. Extending the derivations by Cabré et al. (2007); White et al. (2009); Ross et al. (2011a), we find for the element of the covariance between two galaxy catalogues i, j in two angular bins ϑ, ϑ' :

$$C_{ij}^{\text{TH}}(\vartheta, \vartheta') \equiv \text{Cov}[w^{T_{gi}}(\vartheta), w^{T_{gj}}(\vartheta')] = \sum_l \frac{(2l+1)^2}{(4\pi)^2} \frac{P_l[\cos(\vartheta)] P_l[\cos(\vartheta')]}{\sqrt{f_{\text{sky}}^i f_{\text{sky}}^j}} \text{Cov}(C_l^{T_{gi}}, C_l^{T_{gj}}), \quad (12)$$

where P_l are the Legendre polynomials and the harmonic space covariance is given by

$$\text{Cov}(C_l^{T_{gi}}, C_l^{T_{gj}}) = \frac{1}{2l+1} [C_l^{T_{gi}} C_l^{T_{gj}} + C_l^{TT, \text{tot}} (C_l^{g_{ij}} + \delta_{ij}/n_{s,i})], \quad (13)$$

where δ_{ij} is the Kronecker symbol.

We find that the analytical covariance is comparable with the MC method used in our main analysis; the diagonal elements agree at the $\sim 10\%$ level. The significance of the detections is decreased in this case by $\sim 10\%$: we find for the total combined result $A_{\text{tot}}^{\text{TH}} = 1.42 \pm 0.34$, i.e. $S/N = 4.1$ when using the theory covariance. This is well within the expected fluctuations given the differences in the two procedures.

5.1.2 Absolute versus relative probability

When comparing results from different papers, one must be careful to ensure that they are asking the same questions. As discussed in G08 and S10, there are at least two ways of quantifying the significance of the ISW detection; the method used in G08 and in much of the literature is, *how much is the fit to the data improved by assuming a cross-correlation of the shape predicted by the ISW effect in a Λ CDM model?* By allowing a free amplitude for the expected cross-correlation shape, G08 found $\Delta\chi^2 = 19$ between the best fit amplitude and the hypothesis of no ISW cross-correlation, and the best amplitude suggested by the data was very close to that predicted by Λ CDM.

Another possible approach is simply to ask whether the null hypothesis, that there is *no* cross-correlation, has been ruled out, without assuming any particular alternative model. In G08 it was found that a simple chi-square test is still passed, i.e. the null hypothesis is not rejected by current data, having $\chi_0^2 = 67$ for 73 d.o.f. The absolute chi-square statistic is sensitive to the estimation of the error bars and correlations between measurements; further, ignoring the improvements to the fits when a well-motivated (or even better motivated, given the other evidence) alternative model is considered seems overly pessimistic.

We have checked that the above-mentioned results by G08 remain similar in the current updated version of the data: we obtain for the null hypothesis that $\chi_0^2 = 60.2$, with 78 d.o.f. For the

WMAP7 cosmology this is reduced to $\chi^2_{\text{WMAP}} = 41.6$, and for our best fit amplitude model to $\chi^2_{\text{best}} = 39.7$. At face value, these numbers still do not reject the null hypothesis. Further, the reduced χ^2 is small, potentially indicating that the covariance has been overestimated. Nonetheless, as mentioned above, the most important quantity is the differential $\Delta\chi^2$, which indicates a strong preference for the Λ CDM paradigm.

We have also recalculated these results using the TH covariance: in this case we obtain results more in line with the expectations for this number of degrees of freedom: for the null hypothesis $\chi^2_0 = 88.4$, which is weakly disfavoured at the 80% level, and $\chi^2_{\text{WMAP}} = 72.9$. While the difference between TH and MC results is unclear, it is important to notice that the differential between the null hypothesis and Λ CDM remain consistent at $15 < \Delta\chi^2 < 20$.

5.1.3 Predictions for best possible ISW measurement

While sometimes cited as a skeptical paper regarding the ISW, the work by Francis & Peacock (2010b) (FP) does not contradict the earlier ISW literature, though it is perhaps overly conservative in its discussion of the prospects for measuring the ISW signal. In their study, they repeat the 2MASS cross-correlation measurement, taking advantage of improved redshift determination of the 2MASS sources to divide the sample in three redshift bins. Their measurement of the CMB cross-correlation shows no detection, but a weak preference for Λ CDM models compared to no ISW signal. This is consistent with most measurements using 2MASS (Rassat et al. 2007; Ho et al. 2008; Giannantonio et al. 2008b) apart from the initial claims by Afshordi et al. (2004). The lack of detection is not surprising because most 2MASS objects are at lower redshifts compared to where the ISW signal is expected to arise.

However, FP continue to say that the ISW signal might avoid detection in 10% of cases, even given “the best possible data”. This seems more pessimistic than the well-known limits on the ISW detection, where the maximum signal-to-noise is capped at $(S/N) < 8 - 10$ for the fiducial Λ CDM cosmology. However, this depends significantly on FP’s definition of “the best possible data”; their 10% prediction assumes a maximum redshift of $z_{\text{max}} = 0.7$, which effectively cuts out 50% of the ISW signal, and so is not the best possible data from an ISW perspective. Including the whole redshift range reduces the number of cases where the ISW effect is not detected by an order of magnitude. It is also worth mentioning that FP calculate the fraction of cases where there exists strong evidence for the detection as $\Delta\chi^2 \geq 5$; however, any positive value of $\Delta\chi^2$ would indicate a preference for Λ CDM models compared to no correlation.

In a companion paper, Francis & Peacock (2010a) calculate the expected ISW signal based on the observed 2MASS data. This will be a useful technique when applied to surveys probing redshifts where the ISW is most sensitive, as it provides a template to search for in the CMB maps. For the case of 2MASS, where no cross-correlation is seen, it predicts a larger ISW signal than is seen in WMAP, which is quite surprising given the 2MASS could only be sensitive to a fraction of the ISW present. However, this is primarily due to the much larger than expected power in their highest redshift bin, which has significantly lower galaxy density and is more prone to contamination.

5.1.4 Significance and field to field fluctuations

In another paper, López-Corredoira et al. (2010) argue that the significances of ISW detections claimed by many authors are incor-

rect, based on a mis-estimation of what they call “field-to-field fluctuations”. Their analysis yields comparable results for the CMB-galaxy cross-correlation function to previous detections, but their estimates for the noise in the measurements are significantly higher. G08 and independently Cabré et al. (2006) estimate these uncertainties around $\approx 0.2 \mu\text{K}$, while López-Corredoira et al. (2010) have ≈ 0.35 to $0.6 \mu\text{K}$.

The reason for this disagreement is unclear. The origin and calculation methods appear to be comparable to previous approaches; they briefly discuss jack-knife, Monte Carlo, rotations and analytic methods. For jack-knives, they find comparable answers to previous approaches, while their application of other methods yields much higher noise estimates and they argue that these are more appropriate. We have discussed above the rotation tests, which we view as consistent with other estimates. Analytic approaches have been performed before (Cabré et al. 2007), yielding results comparable to the other covariance methods, and the explanation of the lack of agreement with López-Corredoira et al. (2010) is not clear.

Their differences in the Monte Carlo and analytic results are perhaps hardest to understand; one issue could be that the number of Monte Carlos that they perform (100) is small compared to what is required for many purposes. Earlier calculations typically have used thousands of simulations, in order to ensure convergence in the off-diagonal covariance and invertibility of the covariance matrix. However, fewer Monte Carlos on its own seems unlikely to be able to explain such a large magnitude difference in the diagonal covariance.

It is worth noting that a subset of the authors of López-Corredoira et al. (2010) report a similar inconsistency with estimates of the noise in the galaxy auto-correlation, drawing into question previous BAO measurements (Sylos Labini et al. 2009). This suggests the origin for the cross-correlation discrepancy may relate to their estimation of the large-scale structure fluctuations. However, to this point the discrepancy has not been explained satisfactorily in either the ISW or BAO context.

5.1.5 Comparison to the expected ISW signal

Hernández-Monteagudo (2010) presents an analysis critical of the ISW detections with the NVSS data set, one of the most analysed data sets in this context (Boughn & Crittenden 2004; Nolita et al. 2004; Pietrobon et al. 2006; McEwen et al. 2008; Raccañelli et al. 2008; Schiavon et al. 2012). His results are somewhat mixed; on the one hand he confirms the measurement of a signal seen in the cross-correlation, and in the cross power spectrum ($\ell \sim 10 - 25$) at a weaker $2\text{-}\sigma$ level. However, he claims the theoretical signal should be $5\text{-}\sigma$, significantly larger than what is observed. This statement seems at odds with other measurements using this data, which detected the ISW at the expected level but with a lower significance. It is not clear whether the expected amplitude or the inferred errors are responsible for this difference.

Hernández-Monteagudo (2010) also argues that the correlations arise on smaller scales than expected for the ISW. In the $\ell \sim 2 - 10$ region, which he argues should have half of the ISW signal-to-noise, he finds a low value compared to Ho et al. (2008), leading him to speculate that an unknown foreground systematic may be contributing to the observed ISW signal. Unfortunately, the disagreement with Ho et al. (2008) is not explained, and it is difficult to guess its origin given limited details of the methods.

Hernández-Monteagudo’s statement that half the signal should appear at $\ell < 10$ is somewhat misleading. It is $(S/N)^2$

which adds cumulatively in multipole, not S/N , and only a quarter of the $(S/N)^2$ is expected for $\ell < 10$. It is true that this implies S/N reaches half of its full value in this range, but this would be equally true for three other independent ranges of ℓ . Unfortunately, it is problematic to take a low significance result and attempt to break it up into subsets where the expected signal-to-noise is of order one. Estimates of significance are usually taken into account using template fitting techniques which optimally combine the signal on all scales.

5.1.6 Higher statistical significance

Most detections of the ISW are performed using two-point statistics in real or multipole space, and they largely yield comparable answers. One exception to this is the main galaxy survey detection found by Cabré et al. (2006), which found a higher significance detection (total $S/N = 4.7$ with data from SDSS DR4 alone). This discrepancy was traced by G08 to be due to a particular cut made by Cabré et al. (2006) on the data (they were excluding galaxies with a large error on their Petrosian r magnitude); G08 was not able to justify this cut, and to be conservative used the lower significance answer found without it. The dependence of the answer on this particular cut is somewhat worrying and should be folded into the estimate of the systematic uncertainty. However it should be noted that many other similar cuts have been explored, rarely changing the answers significantly.

In addition, other methods have been used to search for the ISW effect, including wavelet and stacking methods, and these have sometimes produced much higher S/N than is seen using the simpler two-point methods. For example, a number of wavelet methods have been used to analyse the NVSS cross-correlation (Vielva et al. 2006; McEwen et al. 2008; Pietrobon et al. 2006) and have sometimes reported higher significances. However, interpretation of the associated significances are not as straightforward, and often they explore a large space of possible wavelet shapes and scales and report the highest significance detection without reference to the expected theoretical dependence. In such cases, there is an *a posteriori* bias in the statistics; if one focuses on the cosmological constraints from all the measurements, these appear to give comparable constraints on cosmological parameters to those from the two-point statistics.

Using a different approach, Granett et al. (2008) have looked for supercluster and supervoids in the SDSS data; stacking the CMB fields associated with these, they have observed a temperature hot spot at the positions of the superclusters and a cold spot associated with the supervoids. The significance claimed is 4.4σ , much higher than seen in the two-point statistics. Were the fluctuations Gaussian, which is expected on these large scales, we would expect that correlating peaks is not as optimal compared to a full two-point measurement, so the higher significance is surprising. However, there is some *a posteriori* bias imparted when choosing how many superclusters to stack, and how large a patch to consider, which may contribute to the higher significance.

In a later paper (Granett et al. 2009) the same authors generated a linear map of the time derivative of the gravitational potential traced by SDSS LRGs, using a Voronoi tessellation technique. While cross-correlations of this map with the CMB reproduced the expected ISW signal at the expected significance level of $\sim 2\sigma$, this ISW map failed to show any signature associated with the superclusters and supervoids of the earlier detection: the mean temperature of the clusters and voids on the ISW map was not significant, as the temperature difference between clusters and voids was

found to be compatible with zero, contradicting the suggestion that the signal is due to the linear ISW effect.

The excess signal found in Granett et al. (2008) has been investigated by various authors (Papai & Szapudi 2010; Papai et al. 2011; Inoue et al. 2010; Nadathur et al. 2012), who have shown that it is not easy to explain. Their tests have explored a range of different choices of density profiles for super-voids and -clusters and focused on the linear ISW theory as the origin; however no consensus has been reached on the level of disagreement. If this excess signal were to be confirmed and shown to be evidence of a higher-than-expected ISW effect, it could suggest a significantly different cosmological model; however, given the novelty of the method and possible new systematics, it is too early to draw any strong conclusions.

5.2 New galaxy data sets

In the coming years, deeper and wider cosmological surveys will improve the significance of ISW measurements (e.g. DES, PanSTARRS, WISE, LSST, Euclid). In the meantime, existing maps can grow to cover more area and/or be re-analysed to obtain new photometric subsamples. For example, S10 explored three new LRG catalogues at low, medium and high redshifts using the existing SDSS imaging data. The low redshift catalogue has a depth of $z = 0.35 \pm 0.06$, similar to the SDSS main galaxy sample, but has relatively few galaxies making Poisson noise a dominant source of error. The medium redshift 2SLAQ sample most closely approximates the MegaZ LRG sample, but with approximately half as many sources; it has a cross-correlation signal consistent with that seen for the MegaZ data as discussed above. The most interesting ISW result of S10 arises from their high redshift catalogue based on photometric LRGs calibrated using a pilot AAOmega survey, extending to $z = 0.7$. While the low redshift catalogues simply see lower ISW significance than seen in other analyses, the high- z AAOmega sample appears inconsistent with any ISW detection at all.

One concern is that the AAOmega sample is pushing the SDSS photometry to its usable limit. For example, AAOmega selected objects with i -band de Vaucouleurs magnitudes $19.8 < i_{\text{dev}} < 20.5$, significantly deeper than the cut used by SDSS-III BOSS (Eisenstein et al. 2011) and studied in Ross et al. (2011b) of approximately $i_{\text{dev}} < 19.9$. Furthermore, S10 selected their high-redshift LRGs using the r_{iz} colour selection technique as defined by Eqs. (2) through (6) of Ross et al. (2008). As shown in Fig. 1 of Ross et al. (2008), these colour cuts can be quite narrow e.g., their priority B objects are defined to be within a colour range of only $0.2 \leq (i - z) \leq 0.6$, which contains most of the $0.6 < z < 0.7$ AAOmega redshifts. At such faint model magnitudes, we can expect large photometric uncertainties on the observed colours, as demonstrated in Table 6 of Scranton et al. (2005), who found that at $20 < r < 21$ the true (independent) colour errors for red galaxies (like LRGs) is 0.16 for $(i - z)$, i.e., nearly half the width of the AAOmega colour selection discussed above.

We would expect such photometric errors to cause considerable scatter about these AAOmega colour selection boundaries, preferentially leading to contamination from lower redshift and/or lower luminosity interlopers, as there are many more such galaxies in the SDSS sample. Such contamination would lower the correlation function of the AAOmega sample and, interestingly, AAOmega high-redshift sample does possess the lowest auto-correlation function of all three samples used by S10. Finally, it must be remembered that the photometric cuts are based on only

1270 calibration AAOmega spectra, 587 of which were confirmed to be LRGs (see Ross et al. 2008), and this small sample was then extrapolated to 800,000 galaxies in the high- z photometric sample (~ 2000 additional LRG spectra were used to constrain the AAOmega redshift distribution). For comparison, the MegaZ sample we use herein was calibrated using 13,000 spectra.

Stellar contamination is another acknowledged issue. In S10, this contaminant is estimated to be at the 16% level in the full sample of 800 000 high- z LRGs, and at the 9% level in a smaller cleaner data set, as also confirmed by the analysis of the auto-correlation of the LRGs (Sawangwit et al. 2011). However, even this reduced level is still a significant concern and it could undermine the usefulness of this catalogue for ISW measurements where the expected signal is weak. For comparison, the MegaZ LRG and the SDSS QSO samples used in G08 had both a $< 5\%$ stellar contamination, and there it was found that any higher contamination was degrading the data quality beyond usability for the cross-correlation purposes. While it is expected that a component of random stars ought to have null correlation with the CMB, it may correlate with foregrounds (e.g., Galactic dust and stars both trace the structure of the Galaxy), and it will in any case add to the overall noise level, which is already high in these studies. Further, the auto-correlation of the AAOmega sample is shown by Sawangwit et al. (2011) to decrease significantly on large scales when areas with $A_r > 0.1$ are masked, thus showing that there is a correlation between the AAOmega sample and Galactic foregrounds.

These considerations suggest that the AAOmega sample is not well-suited for an ISW measurement. It is interesting to see whether or not the ISW effect can be detected with this data set, but perhaps it should not be interpreted on the same footing as the other, lower redshift SDSS samples. The lack of a cross-correlation found by S10 should be tracked in the future with better quality data at a similar redshift, but at present it is hard to draw conclusions from it on the general nature of the ISW effect.

5.3 Systematics

5.3.1 The rotations as a test of systematics

The rotation tests were discussed in detail above; however, we emphasise that systematic contamination would most likely appear associated with rotations about a single axis, and we find no axis with any particular signal beyond what is expected from the estimated covariance. In general, if the number of outliers were much higher than expected, this would most likely indicate a mis-estimation of the covariance rather than indicating a particular new systematic contaminant. The observed consistency seems to confirm our estimates of the significance of the ISW effect.

5.3.2 Large-scale power in NVSS and bias modelling

While systematic errors are not pointed to by the rotations, that is not to say there are no causes for concern. One significant issue with NVSS, highlighted by Hernández-Montenegro and in other work (Blake & Wall 2002; Blake et al. 2004), is that the auto-correlation signal of NVSS on large scales is significantly higher than what is expected in Λ CDM models, even when considering only the brightest sources in the NVSS survey. This confirms earlier results on the NVSS (Raccanelli et al. 2008) and may imply the existence of a systematic or incorrect modelling of the redshift distribution. (However, no indications of this have been seen in cross-correlations with other surveys.) It could also be a

true physical effect arising from an evolving or a scale-dependent bias, the latter predicted by models of primordial non-Gaussianity (Dalal et al. 2008; Slosar et al. 2008; Afshordi & Tolley 2008; Matarrese & Verde 2008; Giannantonio & Porciani 2010; Xia et al. 2010, 2011).

Generally, cross-correlations should be more robust to foreground contaminants than auto-correlation measurements; in this case the radio foregrounds would have to be correlated with the CMB. Though not impossible given that radio sources can emit in the CMB frequencies, it would be surprising if large-scale cross-correlations were introduced at precisely the level to match cross-correlations seen in other surveys and what is expected theoretically from Λ CDM models.

To further test the possible consequences of the excess power in the ACFs, we have modelled this as an extra low-redshift contamination in the catalogues. We have found that we can reproduce the observed excess large-scale auto-correlations by adding a low-redshift Gaussian spike in the redshift distribution of the form

$$\tilde{\varphi}(z) = \varphi(z) + A_z \cdot \max[\varphi(z)] \cdot \exp\left[-\frac{(z - \mu_z)^2}{2\sigma_z^2}\right], \quad (14)$$

where $\max[\varphi(z)]$ represents the maximum of the unaltered distribution $\varphi(z)$. Some freedom exists with respect to the parameters of the spike (A_z, μ_z, σ_z), which can also be different for each galaxy catalogue; we have found that a conservative combination for these parameters is $\mu_z = 0.02$, $\sigma_z = 0.01$, and $A_z = (0.5, 1, 0.2, 2, 1, 1)$ for the six galaxy catalogues respectively. This is conservative in the sense that it produces more excess power than it is observed. When using these modified redshift distributions, the CCFs remain unaffected, as the ISW effect has a negligible contribution at these low redshifts; on the other hand, the ACFs increase as expected. When using these spectra for the analytical covariance, we then obtain larger error bars on the cross-correlation due to this extra power. For this particular setting, the total ISW significance drops to $\sim 3.4\sigma$.

This is a rather extreme case, since as the low-redshift spike is present in all catalogues, it will produce high correlations between the catalogues, which are at odds with the observed density-density correlation functions. To consider an intermediate, more realistic case, we study the scenario where low- z spikes are added where there is a clear excess in the auto-correlations: main galaxies, NVSS, and quasars only, with $A_z = (0, 0.8, 0, 1.5, 0, 1)$. In this case we find that the total significance is $S/N = 3.8$.

Therefore we estimate that even in the worst-case scenario, low-redshift contamination can not affect the significance of our measurement by more than 1σ , while in more realistic cases this is significantly reduced, and its effect is therefore compatible with the systematic error of $\pm 0.4\sigma$ which we quote.

6 CONCLUSIONS

In this paper we have updated our compilation of ISW measurements to the latest available data, finding consistent results with previous studies, a mild (1σ) excess signal with respect to the expectations from the Λ CDM model, and an updated overall significance of $S/N = 4.4 \pm 0.4$. We have performed additional tests on the combined ISW measurements, finding no evidence for systematics. In particular, we concluded that the rotation test is not an issue for our data, and appears to be of little significance for the data by S10. We have shown that our correlation data remain robust with the latest WMAP7 release of CMB data, and with the final SDSS

DR8 imaging. We have discussed the impact of several issues and criticisms which have arisen in the literature in the recent years, concluding that most of the issues are not serious problems for the use of the ISW as a cosmological probe. We have publicly released all the data, including the maps and the masks for the LSS catalogues. Improvements in the statistical analysis and cosmological consequences will be presented in a forthcoming paper.

It is clear that, if the Λ CDM model is the true underlying model of cosmology, the significance of the ISW effect will remain lower than some other cosmological probes; however, it represents nonetheless a unique signal which allows us to independently confirm the presence of dark energy through its impact on structure growth and potentially detect deviations in how gravity works to build cosmic structures. Upcoming data from ongoing and future surveys, such as e.g. DES, PanSTARRS, WISE, LSST and Euclid, will be crucial for answering these questions, and to push the significance of the ISW detections close to its theoretical limits.

ACKNOWLEDGEMENTS

We thank Stephen Boughn, Martin Kilbinger, Will Percival, Stefanie Phleps, and Alvise Raccanelli for useful discussions. TG acknowledges support from the Alexander von Humboldt Foundation, and from the Trans-Regional Collaborative Research Centre TRR 33 – “The Dark Universe”. RC, BN and AR gratefully acknowledge financial support from the STFC via the rolling grant ST/I001204/1.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

REFERENCES

- Afshordi N., Loh Y.-S., Strauss M. A., 2004, *Phys. Rev.*, D69, 083524
- Afshordi N., Tolley A. J., 2008, *Phys. Rev.*, D78, 123507
- Aihara H., et al., 2011, *Astrophys. J. Suppl.*, 193, 29
- Amanullah R., et al., 2010, *Astrophys. J.*, 716, 712
- Barreiro R. B., Vielva P., Hernandez-Monteagudo C., Martinez-Gonzalez E., 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 747
- Bertacca D., Raccanelli A., Piattella O. F., Pietrobon D., Bartolo N., et al., 2011, *J. Cosmology Astropart. Phys.*, 1103, 039
- Blake C., Ferreira P. G., Borrill J., 2004, *Mon. Not. R. Astron. Soc.*, 351, 923
- Blake C., Wall J., 2002, *Mon. Not. R. Astron. Soc.*, 329, L37
- Boughn S., Crittenden R., 2004, *Nature*, 427, 45
- Boughn S. P., Crittenden R. G., 2002, *Phys. Rev. Lett.*, 88, 021302
- Cabré A., Fosalba P., Gaztañaga E., Manera M., 2007, *Mon. Not. R. Astron. Soc.*, 381, 1347
- Cabré A., Gaztañaga E., Manera M., Fosalba P., Castander F., 2006, *Mon. Not. R. Astron. Soc.*, 372, L23
- Cai Y.-C., Cole S., Jenkins A., Frenk C. S., 2010, *Mon. Not. R. Astron. Soc.*, 407, 201
- Clifton T., Ferreira P. G., Padilla A., Skordis C., 2012, *Phys. Reports*, 513, 1
- Cooray A., 2002, *Phys. Rev.*, D65, 083518
- Corasaniti P.-S., Giannantonio T., Melchiorri A., 2005, *Phys. Rev.*, D71, 123521
- Crittenden R., 2006, in *External Correlations of the CMB and Cosmology*, Fermilab, 2006
- Crittenden R. G., Turok N., 1996, *Phys. Rev. Lett.*, 76, 575
- Dalal N., Dore O., Huterer D., Shirokov A., 2008, *Phys. Rev.*, D77, 123514
- Daniel S. F., Caldwell R. R., Cooray A., Serra P., Melchiorri A., et al., 2009, *Phys. Rev.*, D80, 023532
- De Zotti G., Massardi M., Negrello M., Wall J., 2010, *Astron. Astrophys. Rev.*, 18, 1
- Dunlop J., Peacock J., 1990, *Mon. Not. R. Astron. Soc.*, 247, 19
- Dunsby P., Goheer N., Osano B., Uzan J.-P., 2010, *J. Cosmology Astropart. Phys.*, 1006, 017
- Dupé F.-X., Rassat A., Starck J.-L., Fadili M. J., 2011, *Astron. Astrophys.*, 534, A51
- Efstathiou G., 2004, *Mon. Not. R. Astron. Soc.*, 349, 603
- Eisenstein D. J., et al., 2005, *Astrophys. J.*, 633, 560
- Eisenstein D. J., et al., 2011, *Astron. J.*, 142, 72
- Fosalba P., Gaztañaga E., 2004, *Mon. Not. Roy. Astron. Soc.*, 350, L37
- Fosalba P., Gaztañaga E., Castander F., 2003, *Astrophys. J.*, 597, L89
- Francis C. L., Peacock J. A., 2010a, *Mon. Not. R. Astron. Soc.*, 406, 14
- Francis C. L., Peacock J. A., 2010b, *Mon. Not. R. Astron. Soc.*, 406, 2
- Frieman J., Turner M., Huterer D., 2008, *Ann. Rev. Astron. Astrophys.*, 46, 385
- Frommert M., Enßlin T. A., 2009, *Mon. Not. R. Astron. Soc.*, 395, 1837
- Frommert M., Enßlin T. A., Kitaura F. S., 2008, *Mon. Not. R. Astron. Soc.*, 391, 1315
- Gaztañaga E., Manera M., Multamäki T., 2006, *Mon. Not. R. Astron. Soc.*, 365, 171
- Giannantonio T., Crittenden R., 2007, *Mon. Not. R. Astron. Soc.*, 381, 819
- Giannantonio T., Martinelli M., Silvestri A., Melchiorri A., 2010, *J. Cosmology Astropart. Phys.*, 4, 30
- Giannantonio T., Porciani C., 2010, *Phys. Rev. D*, 81, 063530
- Giannantonio T., Song Y.-S., Koyama K., 2008a, *Phys. Rev.*, D78, 044017
- Giannantonio T., et al., 2006, *Phys. Rev.*, D74, 063520
- Giannantonio T., et al., 2008b, *Phys. Rev.*, D77, 123520
- Gorski K. M., et al., 2005, *Astrophys. J.*, 622, 759
- Goto T., Szapudi I., Granett B. R., 2012, *Mon. Not. R. Astron. Soc.*, 422, L77
- Granett B. R., Neyrinck M. C., Szapudi I., 2008, *Astro-*

- phys. J. Lett., 683, L99
- Granett B. R., Neyrinck M. C., Szapudi I., 2009, *Astrophys. J.*, 701, 414
- Hannestad S., Mirizzi A., Raffelt G. G., Wong Y. Y. Y., 2010, *J. Cosmology Astropart. Phys.*, 8, 1
- Hartlap J., Simon P., Schneider P., 2007, *Astron. Astrophys.*, 464, 399
- Hernández-Monteagudo C., 2008, *Astron. Astrophys.*, 490, 15
- Hernández-Monteagudo C., 2010, *Astron. Astrophys.*, 520, A101
- Hirata C. M., Padmanabhan N., Seljak U., Schlegel D., Brinkmann J., 2004, *Phys. Rev.*, D70, 103501
- Ho S., Hirata C., Padmanabhan N., Seljak U., Bahcall N., 2008, *Phys. Rev.*, D78, 043519
- Ichikawa K., Sekiguchi T., Takahashi T., 2008, *Phys. Rev.*, D78, 083526
- Inoue K. T., Sakai N., Tomita K., 2010, *Astrophys. J.*, 724, 12
- Jarosik N., Bennett C., Dunkley J., Gold B., Greason M., et al., 2011, *Astrophys. J. Suppl.*, 192, 14
- Komatsu E., et al., 2011, *Astrophys. J. Suppl.*, 192, 18
- Lampeitl H., et al., 2010a, *Mon. Not. R. Astron. Soc.*, 401, 2331
- Lampeitl H., et al., 2010b, *Astrophys. J.*, 722, 566
- Larson D., Dunkley J., Hinshaw G., Komatsu E., Nolte M., et al., 2011, *Astrophys. J. Suppl.*, 192, 16
- Lewis A., Challinor A., Lasenby A., 2000, *Astrophys. J.*, 538, 473
- Liu G.-C., Ng K.-W., Pen U.-L., 2011, *Phys. Rev.*, D83, 063001
- Lombriser L., 2011, *Phys. Rev.*, D83, 063519
- Lombriser L., Hu W., Fang W., Seljak U., 2009, *Phys. Rev.*, D80, 063536
- Lombriser L., Slosar A., Seljak U., Hu W., 2012, *Phys. Rev. D*, 85, 124038
- López-Corredoira M., Sylos Labini F., Betancort-Rijo J., 2010, *Astron. Astrophys.*, 513, A3
- Matarrese S., Verde L., 2008, *Astrophys. J. Lett.*, 677, L77
- McEwen J. D., Vielva P., Hobson M. P., Martínez-González E., Lasenby A. N., 2007, *Mon. Not. R. Astron. Soc.*, 376, 1211
- McEwen J. D., Wiaux Y., Hobson M. P., Vanderghenst P., Lasenby A. N., 2008, *Mon. Not. R. Astron. Soc.*, 384, 1289
- Nadathur S., Hotchkiss S., Sarkar S., 2012, *J. Cosmology Astropart. Phys.*, 6, 42
- Nolta M. R., et al., 2004, *Astrophys. J.*, 608, 10
- Padmanabhan N., Seljak U., Pen U., 2003, *New Astron.*, 8, 581
- Padmanabhan N., et al., 2005, *Phys. Rev.*, D72, 043525
- Papai P., Szapudi I., 2010, *Astrophys. J.*, 725, 2078
- Papai P., Szapudi I., Granett B. R., 2011, *Astrophys. J.*, 732, 27
- Percival W. J., et al., 2010, *Mon. Not. R. Astron. Soc.*, 401, 2148
- Pietrobon D., Balbi A., Marinucci D., 2006, *Phys. Rev.*, D74, 043524
- Pogosian L., Corasaniti P. S., Stephan-Otto C., Crittenden R., Nichol R., 2005, *Phys. Rev.*, D72, 103519
- Raccanelli A., et al., 2008, *Mon. Not. Roy. Astron. Soc.*, 386, 2161
- Rassat A., Land K., Lahav O., Abdalla F. B., 2007, *Mon. Not. Roy. Astron. Soc.*, 377, 1085
- Rees M., Sciama D., 1968, *Nature*, 217, 511
- Richards G. T., Myers A. D., Gray A. G., Riegel R. N., Nichol R. C., et al., 2009, *Astrophys. J. Suppl.*, 180, 67
- Ross A. J., Percival W. J., Crocce M., Cabré A., Gaztañaga E., 2011a, *Mon. Not. R. Astron. Soc.*, 415, 2193
- Ross A. J., et al., 2011b, *Mon. Not. R. Astron. Soc.*, 417, 1350
- Ross N. P., et al., 2008, *Mon. Not. Roy. Astron. Soc.*, 387, 1323
- Rozo E., Wechsler R. H., Rykoff E. S., Annis J. T., Becker M. R., et al., 2010, *Astrophys. J.*, 708, 645
- Sachs R. K., Wolfe A. M., 1967, *Astrophys. J.*, 147, 73
- Sawangwit U., Shanks T., Abdalla F. B., Cannon R. D., Croom S. M., Edge A. C., Ross N. P., Wake D. A., 2011, *Mon. Not. R. Astron. Soc.*, 416, 3033
- Sawangwit U., Shanks T., Cannon R. D., Croom S. M., Ross N. P., Wake D. A., 2010, *Mon. Not. R. Astron. Soc.*, 402, 2228
- Schäfer B. M., Douspis M., Aghanim N., 2009, *Mon. Not. R. Astron. Soc.*, 397, 925
- Schäfer B. M., Kalovidouris A. F., Heisenberg L., 2011, *Mon. Not. R. Astron. Soc.*, 416, 1302
- Schiavon F., Finelli F., Gruppiso A., Marcos-Caballero A., Vielva P., Crittenden R. G., Barreiro R. B., Martínez-González E., 2012, *ArXiv e-prints*: 1203.3277
- Scranton R., Connolly A. J., Szalay A. S., Lupton R. H., Johnston D., Budavari T., Brinkman J., Fukugita M., 2005, *ArXiv e-prints*: astro-ph/0508564
- Scranton R., et al., 2003, *ArXiv e-prints*: astro-ph/0307335
- Serra P., Cooray A., Holz D. E., Melchiorri A., Pandolfi S., et al., 2009, *Phys. Rev.*, D80, 121302
- Slosar A., Hirata C., Seljak U., Ho S., Padmanabhan N., 2008, *J. Cosmology Astropart. Phys.*, 0808, 031
- Smail I., Ellis R. S., Fitchett M. J., 1994, *Mon. Not. R. Astron. Soc.*, 270, 245
- Smith R. E., Hernandez-Monteagudo C., Seljak U., 2009, *Phys. Rev.*, D80, 063528
- Smith R. E., Scoccimarro R., Sheth R. K., 2007, *Phys. Rev.*, D75, 063512
- Sylos Labini F., Vasilyev N. L., Baryshev Y. V., López-Corredoira M., 2009, *Astron. Astrophys.*, 505, 981
- Thomas S. A., Abdalla F. B., Lahav O., 2011a, *Physical Review Letters*, 106, 241301
- Thomas S. A., Abdalla F. B., Lahav O., 2011b, *Mon. Not. R. Astron. Soc.*, 412, 1669
- Väliiviita J., Giannantonio T., 2009, *Phys. Rev. D*, 80, 123516
- Väliiviita J., Maartens R., Majerotto E., 2010, *Mon. Not. R. Astron. Soc.*, 402, 2355
- Vielva P., Martínez-González E., Tucci M., 2006, *Mon. Not. R. Astron. Soc.*, 365, 891
- White M., Song Y.-S., Percival W. J., 2009, *Mon. Not. R. Astron. Soc.*, 397, 1348
- Xia J.-Q., Baccigalupi C., Matarrese S., Verde L., Viel M., 2011, *J. Cosmology Astropart. Phys.*, 1108, 033
- Xia J.-Q., Viel M., Baccigalupi C., De Zotti G., Matarrese S., et al., 2010, *Astrophys. J.*, 717, L17
- Xia J.-Q., Viel M., Baccigalupi C., Matarrese S., 2009, *J. Cosmology Astropart. Phys.*, 9, 3
- Yang A., Saslaw W. C., 2011, *Astrophys. J.*, 729, 123
- Zhao G.-B., Giannantonio T., Pogosian L., Silvestri A., Bacon D. J., et al., 2010, *Phys. Rev.*, D81, 103510